

22CT



(12) **EUROPEAN PATENT SPECIFICATION**

(45) Date of publication and mention
of the grant of the patent:
26.09.2001 Bulletin 2001/39

(21) Application number: **98911489.7**

(22) Date of filing: **24.03.1998**

(51) Int Cl.7: **G06F 11/34, G06F 11/32**

(86) International application number:
PCT/US98/04332

(87) International publication number:
WO 98/44418 (08.10.1998 Gazette 1998/40)

(54) **CLIENT-BASED SYSTEM FOR MONITORING THE PERFORMANCE OF APPLICATION PROGRAMS**

KUNDEN-BASIERTES SYSTEM ZUR UEBERWACHUNG DER LEISTUNG EINES ANWENDUNGSPROGRAMMS

SYSTEME CLIENT POUR VERIFICATION DES PERFORMANCES DE PROGRAMMES D'APPLICATION

(84) Designated Contracting States:
AT BE CH DE DK ES FI FR GB GR IE IT LI LU MC NL PT SE

(30) Priority: **01.04.1997 US 831350**

(43) Date of publication of application:
05.07.2000 Bulletin 2000/27

(73) Proprietor: **Landmark Technology Holdings Corporation**
Wilmington, DE 19801 (US)

(72) Inventors:
• **YEE, Andre, H.**
Herndon, VA 20171 (US)
• **FLYNN, Christopher, M.**
Centreville, VA 20120 (US)
• **HARKINS, Russell, W.**
Arlington, VA 22203-1958 (US)

(74) Representative:
DIEHL GLAESER HILTL & PARTNER
Patentanwälte
Augustenstrasse 46
80333 München (DE)

(56) References cited:
EP-A- 0 652 518 WO-A-95/16949
WO-A-96/36918

- **NOBUKO KISHI: "SIMUI: GRAPHICAL USER INTERFACE EVALUATION USING PLAYBACK" PROCEEDINGS OF THE ANNUAL INTERNATIONAL COMPUTER SOFTWARE AND APPLICATIONS CONFERENCE (COMPSAC), CHICAGO, SEPT. 21 - 25, 1992, no. CONF. 16, 21 September 1992, INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, pages 121-127, XP000325751**

Note: Within nine months from the publication of the mention of the grant of the European patent, any person may give notice to the European Patent Office of opposition to the European patent granted. Notice of opposition shall be filed in a written reasoned statement. It shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

Description

Field of the Invention

[0001] The present invention is directed to computer systems, and more particularly to the monitoring of the performance of one or more application programs being executed on a computer system.

Background of the Invention

[0002] As computer systems continue to grow in size, and the capabilities available to end users via such systems continue to expand, there is an ever-increasing demand upon available resources. This is particularly true in the case of computer networks, where a common resource may be shared by a number of client computers. In such a situation, bottlenecks can occur during the execution of application programs that are being run by users on individual clients of the network. The ability to recognize and prevent such bottlenecks from occurring during the execution of application programs has a direct impact on the productivity of the users of the computer system. To this end, therefore, it is desirable to monitor the response time of an application program during execution, to identify potential bottlenecks. By doing so, available resources can be reallocated as necessary to ensure optimum performance.

[0003] In general, one indication of application performance is obtained by measuring the amount of time that is required to respond to a particular request made by or to an application program. To do so, it is therefore necessary to record the times at which the request is generated, and the response is returned. By measuring the latency between the recorded times, a metric is obtained that provides a good indicator of the application's performance. Depending upon the amount of information that is desired, different amounts of monitoring criteria can be established. For example, if a measure of gross latency is sufficient to monitor performance, it may only be necessary to record the time at which the request is initially generated and the time at which the ultimate response is returned. Alternatively, if more information is desired, a time can be recorded for each individual task that is carried out in responding to the request, such as opening windows, performing calculations, etc.

[0004] In the past, the recording of the time at which various events occur during the execution of an application program was accomplished by incorporating application programming interfaces (APIs) into the program whose performance was to be monitored. In essence, these programming interfaces provide a form of instrumentation that permit individual events in the execution of an application to be identified, so that the time of their occurrence can be recorded. One particular advantage of this approach is the fact that any amount of desired detail, at any given level of operation of the pro-

gram, can be obtained.

[0005] However, there are various limitations associated with this "instrumented" approach. One of these is the fact that it is highly labor intensive, since it requires the program to be rewritten to incorporate the APIs. To do so, of course, requires access to the source code for the program. As a result, off-the-shelf programs that are typically sold in "shrink-wrapped" form cannot be adapted by the user to monitor their performance.

[0006] Recently, a standard set of APIs has been proposed to provide greater uniformity in performance monitoring. See "Application Response Measurement API Guide", May 1996. While these standard APIs can be readily incorporated into newly developed programs, they cannot be used with pre-existing programs unless the programs are rewritten to incorporate them. Again, therefore, significant effort must be expended to be able to monitor such programs.

[0007] A different, but related, approach is disclosed in U.S. Patent No. 5,485,574. In the system of this patent, a facility is provided in the kernel of the computer's operating system to count instructions or calls to sections of program code. While this approach avoids the need to incorporate APIs into individual application programs, it still requires that a program, in this case the kernel of the computer's operating system, be modified to provide the necessary instrumentation. In addition, results are obtained at the kernel level of the operating system, rather than at the application or user level.

[0008] As an alternative to the instrumented approach that requires customization of an individual program, another technique for monitoring the performance of a system relies upon network-based communications. Typically, communications between an application program running on a client station and a network server are transmitted via individual data packets. All of the packets pertaining to a given application program are transmitted through a logical port associated with the server. In the network-based approach to monitoring, all packets which pass through a given port are opened, and examined, to determine the types of actions to which they pertain. If a packet contains data relating to a task of interest, the time at which that packet passes through the port is recorded. The advantage of this approach is that it does not require any modification of the application program, and theoretically is applicable to all available programs. However, the results provided by this approach are not completely accurate. In particular, the examination of the data packets takes place along the communication path between the client station and the network server. As such, the times which are recorded in association with each packet only reflect the instance at which the packet passes through the designated port. They do not include additional processing time that may be encountered by the packet after it passes through the port but before the final result is delivered to the requesting application. As another consideration, the examination of the packets is typically carried out by a ma-

chine that is separate from the client station, and hence this approach presents increased hardware expenses. Furthermore, in order to examine the individual packets, the network interface card must operate in a promiscuous mode to permit the port traffic to be examined. Since this mode of operation permits external entities to access the packets, security is compromised.

[0009] It is desirable, therefore, to provide a non-instrumented monitor which does not need to be incorporated into specific programs, and thereby permits the performance of any program to be monitored without modification thereof. Further in this regard, it is desirable to provide such a system which is client-based, so that it provides a true measure of the response time associated with a given task.

Summary of the Invention

[0010] In accordance with the present invention, as defined in detail in the appended independent claims, a system for monitoring the response times of application programs detects events of interest by examining communications between an application program and the computer's operating system. A configuration module permits a user to identify specific events that occur during the operation of application programs, by presenting a sequence of messages that are passed between the application programs and the operating system. Preferably, the messages are described by means of a high-level language which employs readily comprehensible terms that avoid the need for a detailed understanding of the application program itself. The user can then select specific events of interest to be recorded for monitoring purposes. Thereafter, as the application programs execute, the events of interest are recorded. The latency between recorded events represents the response times of the application programs. These response times can be stored in a file for the generation of a report pertaining to application performance.

[0011] Further features of the invention, and the advantages offered thereby, are described hereinafter with reference to an embodiment of the invention illustrated in the accompanying drawings.

Brief Description of the Drawings

[0012]

Figure 1 is a block diagram of a networked computer system in which the present invention can be implemented;

Figure 2 is a flow chart of the steps that are carried out in the servicing of a request;

Figure 3 is a schematic block diagram depicting the exchange of messages between an operating system and client applications;

Figure 4 is a block diagram of the architecture of an application monitoring system;

Figure 5 is a functional block diagram of the operation of the configuration manager module; and Figures 6-8 illustrate display panels that are presented to the user during the definition of a transaction to be monitored.

Detailed Description

[0013] To facilitate an understanding of the present invention, it is described hereinafter with reference to a particular embodiment that is implemented in the context of a computer network that is designed to operate with client applications that run on the Windows® operating system. It will be appreciated, however, that the practical applications of the invention are not limited to this particular implementation. Rather, the principles which underline the invention are applicable to stand-alone computers as well as computer networks, and can be used with a variety of different operating systems.

[0014] A typical network of a type in which the present invention might be employed is illustrated in block diagram form in Figure 1. The network includes a central server 10 that services requests promulgated by various client application programs 12 running on individual nodes 14 of the network. For example, the server may include a database 16 which can be accessed by one or more of the client applications. A typical request from the client applications may be to retrieve certain data from the database, and return it to the requesting application, so that it can be displayed to the user at a given node in a desirable format.

[0015] As an example, a user at one of the network nodes may issue a command to retrieve all records from the database which match a given search criterion, and sort the retrieved records in date order. The steps that are involved in the servicing of this request are illustrated in Figure 2. Referring thereto, the request is initiated when the user generates the instruction to perform the scan, for example by clicking upon an "OK" button in a window, or the like. The command associated with the clicking of this button, including the search parameters established by the user, are transmitted from the user's node to the network server. At the server, the database is scanned, and the records matching the search criteria are retrieved. Thereafter, the appropriate calculations are performed on the retrieved records, to sort them in the requested order. Once the sorting has been completed, the results are returned to the client application at the requesting node. At the node, a suitable window is opened and the results are displayed in the window. The total time between the initial generation of the request and the ultimate display of the results that are responsive to that request is a measure of the latency associated with the servicing of the request. By measuring the length of such latencies, it is possible to obtain a metric of the application response time.

[0016] In accordance with the present invention, this latency is measured by detecting preselected events

that are associated with actions of interest. Typically, operating systems rely upon the interchange of information between the client application and the operating system itself. For example, when the user moves a cursor to a specific position on a display screen and clicks a mouse button, the operating system must inform the client application of the location of the cursor and the fact that the mouse button was clicked at this location. Depending upon the command associated with the clicking of the button, the client application may, in turn, inform the operating system of an action to be carried out, e.g., transmit a request to the network server. The present invention monitors application response times by observing these types of communications between an application program and the computer's operating system. [0017] In the Windows operating system, for example, the exchange of information between the operating system and a client application is carried out by means of "messages." Referring to Figure 3, when an application program is running on a computer, the computer's operating system creates a message queue in the memory space which is allocated to that program. In the example of Figure 3, there are two application programs currently being executed, and hence two message queues have been created in their respective memory spaces. Whenever there is information to be passed to an application program, for instance keystrokes entered by a user or a mouse button click, the operating system places this information in the message queue for that application. The application program periodically sends a command to the operating system to get its messages. In response, the operating system passes the contents of the message queue to the application. Conversely, whenever the application program has information to pass to the operating system, e.g., data to be transmitted to the network server 10, the application program may send a command to the operating system via a message. In response, the operating system reads the message, determines the action to be taken, and proceeds accordingly.

[0018] In accordance with the present invention, the application response time is monitored by examining the contents of the communications between the operating system and an application program, to detect certain events of interest. The latency between the selected events is recorded as the measure of response times.

[0019] The observance of communications to monitor response times is carried out by a program known as an applications agent. Referring again to Figure 1, an applications agent 18 might reside on each node 14 of the network having a client application 12 whose response times are to be monitored. The applications agent is external to the application program itself that is to be monitored, and runs in the background in a manner that is transparent to the user. The architecture of the overall application monitoring system is illustrated in Figure 4. The main components of the monitoring system are the application agent 18, a central dispatcher

28 and a configuration manager 30. The configuration manager 30 enables the user to define the events that are to be monitored during the execution of an application program. For this purpose, the configuration manager includes various display panels 32, by which the user establishes the parameters for the monitoring of an application. The operation of the configuration manager will be described with reference to a functional block diagram shown in Figure 5 and exemplary display panels illustrated in Figures 6-8.

[0020] In general, the configuration manager operates by monitoring an application program while it performs a task of interest, recording events that occur in connection with the task, and enabling the user to identify selected events as the measuring criteria for application response times. To define a transaction to be monitored, the user first launches the application, or applications, to be monitored, if they are not already running on the computer. The user then issues appropriate commands, for example through one or more interface windows (not shown) to indicate that a transaction is to be defined.

[0021] In response to such a command, a Transaction Recording dialog box can be displayed, for example of the type illustrated in Figure 6. This dialog box provides the user with two accessible controls for recording messages relating to a transaction, namely "start" and "stop". As an alternative to a dialog box, it is also possible to employ commands from a drop-down menu, or predefined keystrokes. When the applications running on the computer reach a point at which the user desires to define a task to be monitored, the user starts transaction recording, by clicking upon the "Start" button in the dialog box, or pressing an appropriate key on the computer's keyboard. This action activates a message recording function within the configuration manager 30. Thereafter, the user can perform any tasks that are necessary to the operation to be monitored. For example, if the user is interested in the latency associated with the servicing of a request to retrieve data from a database, the user generates the necessary commands to cause an application to carry out this function. Once the transaction has been performed, the "Stop" button, or an appropriate keyboard key, is pressed to terminate the recording function.

[0022] Referring again to Figure 4, as the transaction is being recorded, all of the messages that are exchanged between the computer's operating system and currently running application programs are detected and stored by the central dispatcher 28. Once the recording of the transaction has been terminated, all of the stored messages are provided to the configuration manager. In response, the configuration manager displays a Transaction Macro dialog box 38, as schematically illustrated in Figure 5. A more detailed example of such a dialog box is shown in Figure 7. This dialog box presents to the user a sequential display of each message that was generated during the recorded portion of the trans-

action. For ease of understanding, the particular example shown in Figure 7 illustrates messages that are exchanged between the operating system and a single application program. In practice, however, multiple applications may be running on the computer concurrently. In such a case, the dialog window 38 lists all of the various messages that are exchanged between the operating system and the executing programs during the transaction recording period.

[0023] From the listing in the dialog box, the user can select specific messages that identify the beginning and end of the transaction that is to be monitored. This is accomplished in the illustrated Transaction Macro dialog box by clicking a mouse button while the cursor is positioned in a column labelled "Begin/End", adjacent the messages of interest. In the specific example of Figure 7, the user has selected the clicking of a "Scan" button as the beginning of a marked transaction, and the subsequent clicking of an "OK" button to end the transaction. Of course, other techniques for identifying the beginning and ending statements for a transaction can be employed as well.

[0024] In the example of Figure 7, a single event is displayed for the action which initiates the scanning of the database to retrieve certain results. Once the results are returned to the client application, a window is created and the results are displayed within the window. Subsequently, the user clicks an "OK" button to remove the display of those results. Although only three events are depicted in the dialog box of Figure 7 for this activity, in practice numerous messages might be exchanged between the operating system and the application program to perform the required task. Furthermore, the actual terms employed in those messages may not be readily understandable to the average user. To this end, therefore, the central dispatcher 28 preferably employs a high-level language, such as a macro language 36, to translate the specific content of the messages exchanged between the operating system and the application programs into terms that can be readily understood by the average user. Figure 7 illustrates examples of such terms, which can be used to identify when a window is created or destroyed, and when a user clicks on a button within a window. Each term has associated parameters that identify the function associated with that term. For example, the "WindowCreated" term is followed by three parameters which identify the name of the application with which the window is associated, the title of the window, and the chain of parent windows within which the window being created is contained. To actually create a window, or to exchange information relating to the clicking of a button, several messages may pass between the application program and the operating system. Preferably, the macro language 36 encapsulates several of these messages, and/or events, into a single term that is displayed in the transaction macro dialog window of Figure 7.

[0025] As alternatives, it is possible to employ other

types of high-level languages to perform this function. For example, a script language, a programming language, or a code generator could be employed to translate low-level messages exchanged between the applications programs and the operating system into higher level terms that are more readily understandable to the user.

[0026] Once the user has selected the beginning and ending message statements that mark a transaction to be recorded, that particular transaction can be given a unique name, for subsequent use. At this point, the definition of the transaction to be monitored is complete. Once the transaction has been defined, it can be distributed to agents 18 on the nodes in the computer system where the transaction is to be monitored, as depicted in Figure 5. To this end, the configuration manager 30 preferably provides the user with a view of all agents running on the network, so that the defined transaction can be distributed to selected agents.

[0027] The primary function of a defined transaction is to measure the time that elapses between the beginning and ending events identified by the user. Referring again to Figure 4, when a transaction is to be monitored, the beginning and ending messages are supplied from the agent 18 to the central dispatcher 28. The dispatcher initiates a message watch function 34 which monitors each message that appears in the message queue for an application of interest. In the Windows operating system, for example, the contents of a message queue can be observed without affecting the message by means of a suitable message hook, which examines each message before it is passed to the application. Whenever a designated beginning or ending message appears in the queue, the central dispatcher informs the agent 18, which records the time at which the message appeared. For each recorded pair of start and stop times, the agent 18 computes the latency between the detected events, or messages. The calculated latencies are then stored in a file, for the subsequent generation of a report.

[0028] In the preceding example, only one pair of beginning and ending messages was identified to monitor response times. To obtain a greater amount of detail about the operation of a program, it may be desirable to establish nested sets of start and stop times. For example, within the latency associated with the performance of a general task, the user may wish to also measure the latencies of more specific tasks that are carried out in the implementation of the general task. In a preferred embodiment of the invention, therefore, the user can designate multiple start and stop times for a given transaction to be monitored. During the transaction, the latencies associated with each pair of associated start and stop times are calculated and stored, for subsequent reporting purposes.

[0029] In addition to measuring the latency associated with the performance of tasks, the agent 18 permits other types of network monitoring functions to be carried out. To this end, the user can establish a policy that can

be applied to a defined transaction. In essence, a policy comprises a saved definition of the agents which are to monitor a defined transaction, additional metrics that can be measured in addition to the defined latency period, and specific times at which the monitoring operation is to be carried out. Figure 8 illustrates an example of a dialog box for selecting the characteristic properties of a policy. The specific example illustrated in the figure pertains to various metrics that can be measured during the execution of a program. The right panel 40 in the figure provides a list of measurable parameters for the network system. The user can select these parameters as desired, to be included in the properties of the policy. The selected metrics are illustrated in the left window 42 of the dialog box.

[0030] From the foregoing, therefore, it can be seen that the present invention provides an application monitoring system which is not intrusive, and therefore can be utilized in connection with any type of application program that is executed on a computer. Since the monitoring functions are carried out at the site of the client application itself, rather than intermediate the client and a network server, a true measurement is obtained of the actual latencies that occur in the execution of a program. Furthermore, by utilizing a macro language that translates communications between the application program and operating system into terms that are readily understandable by a user, the ability to configure the monitoring system is greatly enhanced.

[0031] With reference to Figure 4, it can be seen that the monitoring system of the present invention is comprised of a number of components. In a practical implementation of the invention, these components can be distributed in a variety of manners. When used on a stand-alone computer, of course, all of the components may reside on a single computer. However, it is also possible to first define the transaction to be monitored with a configuration manager that is running on one computer, and then distribute that transaction, and any associated policy, to agents on other stand-alone computers which do not contain the configuration manager themselves. In a network environment, it is possible to have an agent running on each network node, as illustrated in the embodiment of Figure 1. However, such an arrangement is not necessary. Rather, the agents can be located at a central location, and from there instruct the central dispatcher, which might also be located at the central location or on the individual nodes, to initiate the monitoring procedures by establishing the watch functions in the memory space of the application programs to be monitored. Thus, as long as the various modules can communicate with each other as necessary, they can be located anywhere within the networked system.

[0032] It will be appreciated by those of ordinary skill in the art that the application can be embodied in other specific forms. For example, although the disclosed embodiment of the invention has been described in the context of a computer which employs the Windows operat-

ing system, the practical applications of the invention are not limited thereto. Rather, the invention can be employed in any type of system in which it is possible to externally detect specific events that are associated with the execution of an application program. The presently disclosed embodiments are therefore considered in all respects to be illustrative, and not restrictive. The scope of the invention is indicated by the appended claims, rather than the foregoing description, and all changes that come within the meaning and range of equivalents thereof are intended to be embraced therein.

Claims

1. A method for monitoring the response times of application programs running on a computer, comprising the steps of:

observing communications between an application program (12) and an operating system program running on the computer;
detecting the times at which predefined communications are exchanged between the application program (12) and the operating system program;
determining the latency associated with an application response in accordance with the difference in detected times; and
storing the determined latency.

2. The method of claim 1 wherein said predefined communications are selected by a user from a sequence of communications that are exchanged between the application program and the operating system program.

3. The method of claim 2 wherein said sequence of communications is provided to the user by operating the application program in a normal manner, recording communications that are exchanged during such operation, terminating the recording of communications and displaying the recorded communications upon termination of the recording process.

4. An applications agent (14) for monitoring the response times of application programs running on a computer, comprising:

means for observing communications between an application program (12) and an operating system program running on the computer;
means for detecting the times at which said communications are exchanged between the application program (12) and the operating system program;
means for calculating the latency associated with the difference in detected times; and

means for storing the calculated latency.

5. The agent of claim 4 further including a configuration manager for enabling a user to select predefined communications from a sequence of communications that are exchanged between the application program and the operating system program, and wherein said detecting means detects the times at which said predefined communications occur. 5
6. The agent of claim 5 wherein said configuration manager includes means for recording communications that are exchanged during the operation of the application program in a normal manner, means for terminating the recording of communications, means for displaying the recorded communications upon termination of the recording process, and means for enabling the user to designate selected ones of the displayed communications as said predefined communications. 10 15 20
7. The agent of claim 6 further including means for translating communications exchanged between the application program and the operating system program into high-level language terms for display to the user. 25
8. The agent of claim 7 wherein said translating means combines plural communications into a single high-level language term. 30
9. A system for monitoring the response times of application programs (12) being executed on a computer, comprising: 35
 - a configuration manager (30) for presenting to a user a sequential listing of detectable events which occur during the execution of an application program (12) including means for enabling the user to select individual ones of said events to define a transaction to be monitored; 40
 - means for examining information exchanged between the application program (12) and an operating system running on the computer, to detect the occurrence of the events selected by the user; 45
 - means for recording the times at which detected events occur; and
 - means for calculating and storing the elapsed time between the occurrence of selected events, to thereby provide a metric of application response time. 50
10. The system of claim 9 wherein said configuration manager includes means for recording events which occur as the user issues commands during the execution of an application, and displaying said recorded events as said sequential listing. 55

Patentansprüche

1. Verfahren zum Beobachten von Antwortzeiten von Anwendungsprogrammen, die auf einem Computer laufen, umfassend die Schritte:
 - Beobachten von Benachrichtigungen zwischen einem Anwendungsprogramm (12) und einem Betriebssystemprogramm, die auf dem Computer laufen;
 - Ermitteln von Zeiten, zu denen vorbestimmte Benachrichtigungen zwischen dem Anwendungsprogramm (12) und dem Betriebssystemprogramm ausgetauscht werden;
 - Bestimmen einer Wartezeit, die mit einer Anwendungsantwort in Übereinstimmung mit einer Differenz in den ermittelten Zeiten verbunden ist; und
 - Speichern der ermittelten Wartezeit.
2. Verfahren gemäß Anspruch 1, wobei die vorbestimmten Benachrichtigungen durch einen Benutzer aus einer Sequenz von Benachrichtigungen ausgewählt werden, die zwischen dem Anwenderprogramm und dem Betriebssystemprogramm ausgetauscht werden.
3. Verfahren gemäß Anspruch 2, wobei die Sequenz von Benachrichtigungen dem Benutzer durch Anwendung des Anwendungsprogramms auf normale Art und Weise, Aufzeichnung von Benachrichtigungen, die während einer solchen Anwendung ausgetauscht werden, beenden der Aufzeichnung der Benachrichtigungen und Anzeigen der aufgezeichneten Benachrichtigungen nach Beendigung des Aufzeichnungsvorgangs bereitgestellt wird.
4. Anwendungsagent (14) zum Beobachten der Antwortzeiten von Anwendungsprogrammen, die auf einem Computer laufen, umfassend:
 - eine Einrichtung zum Beobachten der Benachrichtigungen zwischen einem Anwendungsprogramm (12) und einem Betriebssystemprogramm, die auf einem Computer laufen;
 - eine Einrichtung zum Ermitteln der Zeiten, zu denen die Benachrichtigungen zwischen dem Anwendungsprogramm (12) und dem Betriebssystemprogramm ausgetauscht werden;
 - eine Einrichtung zum Errechnen einer Wartezeit, die mit einer Differenz in den ermittelten Zeiten verbunden ist; und

eine Einrichtung zum Speichern der errechneten Wartezeit.

5. Agent gemäß Anspruch 4, ferner umfassend einen Konfigurationsmanager, um es einem Benutzer zu ermöglichen, vorbestimmte Benachrichtigungen aus einer Sequenz von Benachrichtigungen, die zwischen dem Anwendungsprogramm und dem Betriebssystemprogramm ausgetauscht werden, auszuwählen, und wobei die Ermittlungseinrichtung die Zeiten, zu denen die vorbestimmten Benachrichtigungen stattfinden, ermittelt.

6. Agent gemäß Anspruch 5, wobei der Konfigurationsmanager eine Einrichtung zum Aufzeichnen von Benachrichtigungen, die während der Anwendung des Anwendungsprogramms auf normale Art und Weise ausgetauscht werden, eine Einrichtung zum Beenden der Aufzeichnung von Benachrichtigungen, eine Einrichtung zum Anzeigen der aufgezeichneten Benachrichtigungen nach Beendigung des Aufzeichnungsvorgangs, und eine Einrichtung, die es dem Benutzer ermöglicht, ausgewählte der angezeigten Benachrichtigungen als die vorbestimmten Benachrichtigungen zu kennzeichnen.

7. Agent gemäß Anspruch 6, ferner umfassend eine Einrichtung zum Übersetzen von Benachrichtigungen, die zwischen dem Anwendungsprogramm und dem Betriebssystemprogramm ausgetauscht werden, in Hochsprache, um sie dem Benutzer anzuzeigen.

8. Agent gemäß Anspruch 7, wobei die Übersetzungseinrichtung mehrere Benachrichtigungen zu einem einzelnen Hochsprachenausdruck kombiniert.

9. System zum Beobachten der Antwortzeiten von Anwendungsprogrammen (12), die auf einem Computer ausgeführt werden, umfassend:

einen Konfigurationsmanager (30), um einem Benutzer eine sequentielle Auflistung von ermittelbaren Ereignisse anzuzeigen, die sich während der Ausführung eines Anwendungsprogramms (12) ereignen, umfassend eine Einrichtung, die es dem Benutzer ermöglicht, individuelle der Ereignisse auszuwählen, um eine zu überwachende Durchführung zu definieren;

eine Einrichtung zum Untersuchen von Informationen, die zwischen dem Anwendungsprogramm (12) und einem Betriebssystem, die auf dem Computer laufen, ausgetauscht werden, um das Eintreten der von dem Benutzer ausgewählten Ereignisse zu ermitteln;

eine Einrichtung zum Aufzeichnen der Zeiten,

zu denen sich ermittelte Ereignisse ereignen; und

eine Einrichtung zum Errechnen und Speichern der verstrichenen Zeit zwischen dem Eintreten von ausgewählten Ereignissen, um dadurch eine Metrik der Anwendungsantwortzeit bereitzustellen.

10. System gemäß Anspruch 9, wobei der Konfigurationsmanager eine Einrichtung zum Aufzeichnen von Ereignissen umfasst, die eintreten, wenn der Benutzer Befehle während der Ausführung einer Anwendung erteilt, und zum Anzeigen der aufgezeichneten Ereignisse als die sequentielle Auflistung.

Revendications

1. Procédé pour le contrôle des temps de réponse de programmes d'application se déroulant dans un ordinateur, comprenant les étapes suivantes :

l'observation des communications entre un programme d'application (12) et un programme du système d'exploitation se déroulant dans l'ordinateur ;

la détection des moments auxquels des communications prédéfinies sont échangées entre le programme d'application (12) et le programme du système d'exploitation ;

la détermination de la latence associée à une réponse d'application selon la différence entre les moments détectés ; et

le stockage de la latence déterminée.

2. Procédé selon la revendication 1, dans lequel lesdites communications prédéfinies sont sélectionnées par un utilisateur à partir d'une séquence de communications qui sont échangées entre le programme d'application et le programme du système d'exploitation.

3. Procédé selon la revendication 2, dans lequel ladite séquence de communications est fournie à l'utilisateur en mettant en oeuvre le programme d'application de manière normale, en enregistrant les communications qui sont échangées pendant cette mise en oeuvre, en mettant fin à l'enregistrement des communications et en affichant les communications enregistrées lorsqu'il est mis fin au procédé d'enregistrement.

4. Agent d'application (14) pour contrôler les temps de réponse de programmes d'application se déroulant dans un ordinateur, comprenant :

des moyens pour observer des communica-

- tions entre un programme d'application (12) et un programme du système d'exploitation se déroulant dans l'ordinateur ;
des moyens pour détecter les moments auxquels lesdites communications sont échangées entre le programme d'application (12) et le programme du système d'exploitation ;
des moyens pour calculer la latence associée à la différence entre les moments détectés ; et
des moyens pour stocker la latence calculée.
5. Agent selon la revendication 4, comprenant en outre un gestionnaire de configuration pour permettre à un utilisateur de sélectionner des communications prédéfinies à partir d'une séquence de communications qui sont échangées entre le programme d'application et le programme du système d'exploitation, et dans lequel lesdits moyens de détection détectent les moments auxquels se produisent lesdites communications prédéfinies.
6. Agent selon la revendication 5, dans lequel ledit gestionnaire de configuration comprend des moyens pour enregistrer des communications qui sont échangées pendant la mise en oeuvre du programme d'application de manière normale, des moyens pour mettre fin à l'enregistrement des communications, des moyens pour afficher les communications enregistrées lorsqu'il est mis fin au procédé d'enregistrement, et des moyens pour permettre à l'utilisateur de désigner certaines communications affichées sélectionnées comme communications prédéfinies.
7. Agent selon la revendication 6, comprenant en outre des moyens pour traduire des communications échangées entre le programme d'application et le programme du système d'exploitation en termes de langage évolué à des fins d'affichage pour l'utilisateur.
8. Agent selon la revendication 7, dans lequel lesdits moyens de traduction combinent des communications plurielles en un terme unique de langage évolué.
9. Système de contrôle des temps de réponse de programmes d'application (12) à mettre en oeuvre dans un ordinateur, comprenant :
- un gestionnaire de configuration (30) pour présenter à un utilisateur un listage séquentiel d'événements détectables qui se produisent pendant la mise en oeuvre d'un programme d'application (12) comprenant des moyens pour permettre à l'utilisateur de sélectionner individuellement certains desdits événements pour définir une transaction à contrôler ;
- des moyens pour examiner les informations échangées entre le programme d'application (12) et un programme du système d'exploitation se déroulant dans l'ordinateur, pour détecter la survenue des événements sélectionnés par l'utilisateur ;
des moyens pour enregistrer les moments auxquels se produisent les événements détectés ; et
des moyens pour calculer et stocker le temps écoulé entre la survenue des événements sélectionnés, pour fournir ainsi une mesure du temps de réponse d'application.
10. Système selon la revendication 9, dans lequel ledit gestionnaire de configuration comprend des moyens pour enregistrer des événements qui se produisent lorsque l'utilisateur lance des ordres au cours de l'exécution d'une application, et pour afficher lesdits événements enregistrés en tant que dit listage séquentiel.

FIG. 1


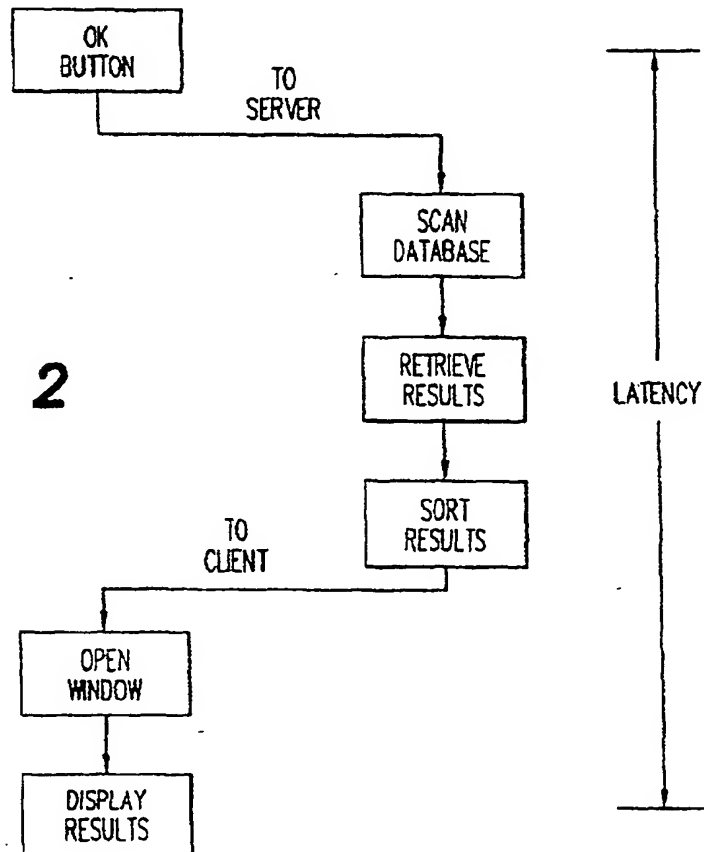


FIG. 2






FIG. 3




FIG. 5




FIG. 4




FIG. 6

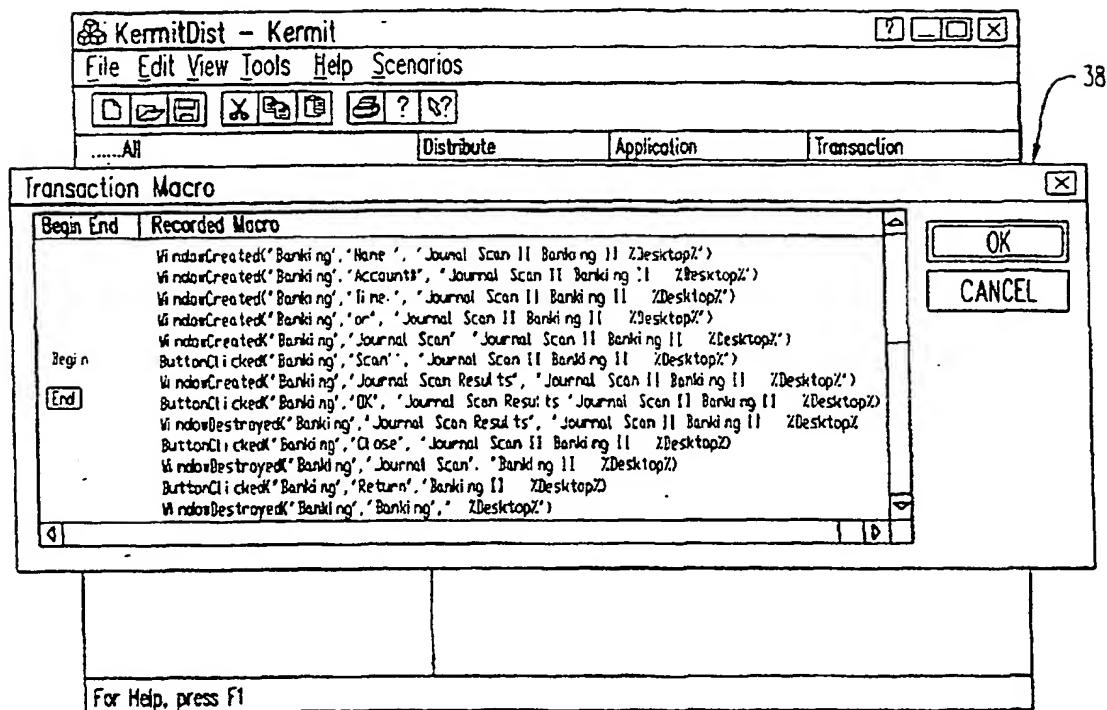


FIG. 7

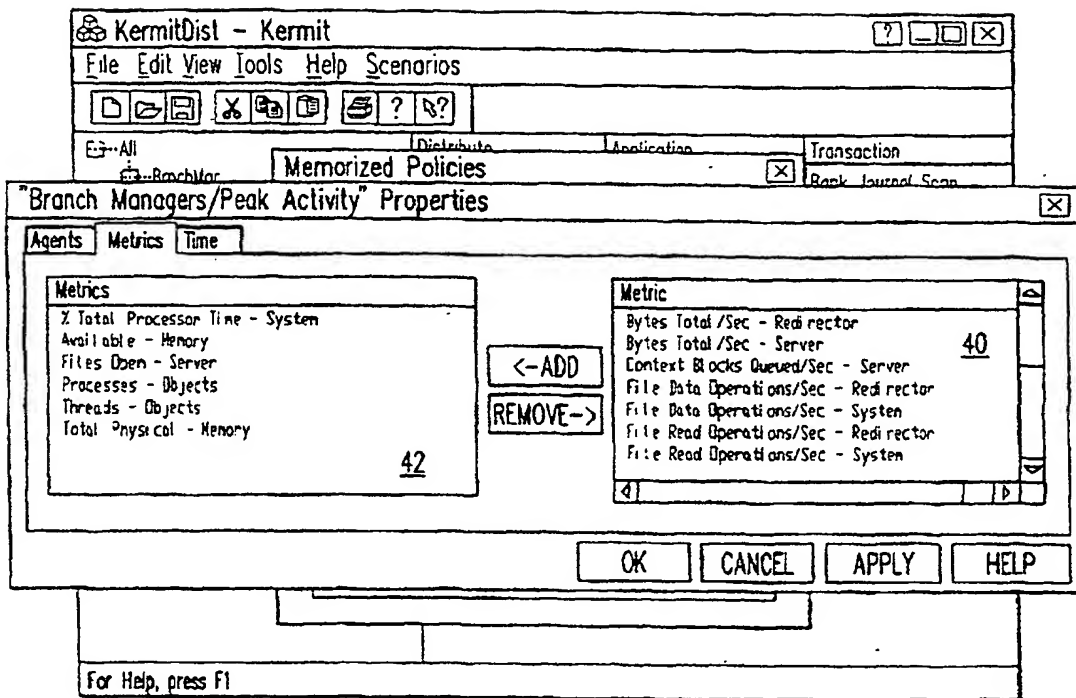


FIG. 8